

APPLICATION
FOR
UNITED STATES LETTERS PATENT

TITLE: NETWORK COMMUNICATION

APPLICANT: YLIAN SAINT-HILAIRE AND BILL STRAHM

CERTIFICATE OF MAILING BY EXPRESS MAIL


Express Mail Label No. EL724384220US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

March 16, 2001

Date of Deposit

Signature



Gil Vargas

Typed or Printed Name of Person Signing Certificate

NETWORK COMMUNICATION

BACKGROUND

This invention relates to network communication.

Mobile devices, such as laptop computers, mobile phones,
5 and handheld digital assistants, for example, communicate on
wireless networks to enable their users to work in multiple
locations and while in transit, such as in homes, airports,
airplanes, and public spaces. Typically, the mobile devices
run applications that access a server on the Internet by
10 communicating using a single "connection."

DESCRIPTION OF DRAWINGS

Figure 1 is a diagram of network connections between a
mobile client and a home agent.

Figure 2 is a schematic of a tunnel state machine.

15 Figure 3 is a flow chart of a process for connecting a
mobile client.

Figure 4 is a schematic of a mobile device.

Figure 5 is a schematic of a stack for providing
information to an application.

20 Figure 6 is a schematic of information processing units
implemented by a mobile client.

Figure 7 is a schematic of information processing units
implemented by a home agent or server.

DETAILED DESCRIPTION

A seamless and robust information exchange protocol for connecting a device to a secure network can include two or more open connections. The connections may be wireless and 5 may be selectively and efficiently used to enhance information transmission speed and reliability.

Referring to the example shown in Figure 1, a mobile (or immobile) device 110 establishes four exemplary connections or tunnels with a home agent 160. Examples of mobile devices 10 include laptop computers, handheld computers (e.g., PalmPilots™, handheld digital assistants, key pads), pagers, mobile telephones (e.g., Wireless Application Protocol (WAP) enabled phones), Moving Picture Experts Group Audio Layer 3 (MP3) players, and game and/or entertainment stations. 15 Examples of immobile devices include desktop computers, large appliances, and factory machines.

The terms "connection" and "tunnel" refer to a route for sending and/or receiving information. Information refers to both data (e.g., text, numeric, Boolean, addresses, graphical 20 content, and the like) and commands (e.g., requests, instructions, queries and the like). The home agent can be a computer system (such as a desktop computer, a server, a server farm, or a mainframe).

Connection 114 is a wireless phone link (e.g., third 25 generation wireless networks (3G), second-and-a-half generation wireless networks (2.5G), General Packet Radio Service (GPRS)), i-mode™ (NTT DoCoMo, Japan), Global System

for Mobile Communications (GSM), Code Division Multiple Access (CDMA), or Time Division Multiple Access (TDMA) / Digital Advanced Mobile Phone Service (D-AMPS) wireless link). This connection can be a modem connection that is relayed to an 5 Internet service provider that connects the mobile client 110 to the Internet 140.

Connection 116 is a wireless connection to a local area network (LAN) 130 (e.g., a wireless LAN such as provided by wireless Ethernet, BlueTooth, the Institute of Electrical and 10 Electronics Engineers (IEEE) 802.11a or 802.11b standard (IEEE std. 802.11-1999, published 1999), or Cellular Digital Packet Data (CDPD) connections). Information from this connection can optionally traverse a Foreign Network Address Translator (NAT) /Proxy 132 to connect from the intranet 130 to the 15 Internet 140.

Connection 118 is a wired network connection, e.g., to a corporate intranet 130. Information from this connection also traverses the Foreign NAT/Proxy 132 to connect to the Internet 140.

20 Connection 120 is a direct connection to the Internet 140, e.g., a Digital Subscriber Line (DSL), Integrated Services Digital Network (ISDN), or cable modem connection. The Internet 140 directs the connections to the home agent 160 that is linked to the Internet by a corporate firewall 152. 25 The home agent 160 can also be linked to a second intranet 150.

The mobile client 110 can be designed to flexibly use a variety of communication protocols (e.g., Transmission Control Protocol (TCP)/Internet Protocol (TCP/IP), User Datagram Protocol (UDP), Hypertext Transfer Protocol (HTTP), Simple Mail Transfer Protocol (SMTP), Post Office Protocol 3 (POP3), or other Internet protocols) for the purpose of communicating to the home agent 160, e.g., on the intranet 150. In some examples, the mobile client 110 is assigned an Internet Protocol (IP) address on the secure intranet 150.

5 Applications running on the mobile client 110 issue network calls using the assigned IP address, typically, using a proxy module on the home agent 160. However, the connections between the mobile client 110 and the home agent 160 may require a different set of addresses. For example, the mobile

10 client 110 may be a guest system having a guest IP address on a second intranet 130 or may be connected to the Internet 140 using a temporary Internet IP address, e.g., through a third party service provider. Addresses, such as the guest IP address and the temporary IP address, are invisible to applications running on the mobile client 110, but are used,

15 e.g., by the network interfaces or operating system, to connect to the home agent 160.

20

The communication connections may use network address translation (NAT). NAT enables machines within a secure intranet to use a set of assigned IP addresses. NAT translates these assigned IP addresses to external IP addresses for communication with devices outside of the

25

intranet. This protocol enhances network security and allows more efficient assignment of IP addresses within an intranet. Generally, the mobile client 110 will establish connections with the home agent through NAT, e.g., Foreign NAT/Proxy, 132.

5 NAT 132 forwards information, e.g., across a public or insecure network, to the home agent 160 that resides securely behind a firewall 152. The home agent 160 processes the information from the mobile client 110, e.g., decompresses and decrypts the information, to allow the mobile client 110 to

10 function securely with the advantages of a direct connection to an intranet or virtual private network (VPN) 150.

Referring to Figure 2, each connection operates as a finite state machine. The connection can be disconnected 210, connecting 220, passive 230, or active 240. In the active state, a connection sends and receives information with the server. An active connection 240 is switched to passive 230 under certain circumstances, for example, if another connection is more efficient, secure, or economical. A user or system administrator may define parameters for such a switch. The system may monitor the active connection 240 and/or the passive connections 230 for one or more of the following parameters: signal strength, transmittal rate, latency, cost of transmittal, and connection integrity.

Monitoring and switching of connections may be performed

25 without the cognizance of the user.

An active connection 240 is switched to the disconnected state 210, for example, in the event of a loss of connection

integrity (e.g. an abrupt break or a security breach). A disconnected connection 210 is switched to connecting 220 when appropriate.

A connecting connection 220 is switched to passive 230
5 once it is established and properly initialized. A passive connection 230 is generally not used to send or receive information except as necessary to monitor or maintain its integrity. In particular, a passive connection 230 is not used to send or receive application information. The passive
10 connection 230 may be used to send instructions for its own activation. For example, if there is an abrupt disconnection of the active connection 240, either or both of the server 160 and the mobile client 110 can issue commands on a passive connection 230 indicating that it should be activated.

15 More than one connection may be active at a given time. Information is divided between such active connections 240, e.g., in proportion to the transfer rate of the particular connection, its latency, its reliability and so forth. The use of multiple connections at one time increases overall
20 information transfer rates.

Referring now also to Figure 3, the mobile client 110 scans 310 available network interfaces checking for ones that permit connection to the home agent 160. The connection process includes initializing communication protocols 320.
25 Typically, the steps included in process 320 are only required for the initial connection 312. Thus, adding additional

passive connections 230 and switching among active connections 240 are rapid and efficient.

To initialize a connection 320, network addressing 322 is initialized. For example, the mobile client 110 receives a 5 Dynamic Host Configuration Protocol (DHCP) address from the home agent 160 or from another server. The mobile client 110 can also detect the presence of a proxy module on the home agent 160. Security protocols are established and authenticated 324 as are compression protocols 326. Examples 10 of security protocols include transport layer security (TLS), secure sockets layers (SSL), and wireless TLS (WTLS). Non-limiting examples of a compression protocol include: "Lossless Compression" (LZS®) and Microcom Networking Protocol-5 (MNP5). Optionally, the mobile client 110 and home agent 160 exchange 15 secondary authentication information, such as authenticated names, digital certificates and cookies.

Once the protocols are initialized for a session, initiation of additional connections (e.g., at least a second connection) to the home agent 160 only requires authentication 20 314 to verify that a new connection is between a mobile client 110 and/or a home agent 160 that are already party to a session. The same network, security, and compression 25 protocols and parameters are used for information exchange as for the initial connection. For example, the common parameters for multiple sessions may include cryptographic keys, network address assignments, and compression formats (e.g., compression dictionaries). Accordingly, in such

examples, all passive 230 and active connections 240 use the same security level. The precise level can be determined by a system administrator or user.

In other examples, the system uses different security protocols and/or compression levels for the different connections. The home agent 160 and mobile device 110 establish different security procedures for each connection, e.g., each passive or active tunnel, depending on the connection type. Thus, a connection that requires firewall traversal, e.g., a mobile phone link to the Internet, requires high level security, whereas a wired Ethernet connection to a LAN does not require such a high level of security.

The mobile client 110 and/or home agent 160 can be programmed to determine if additional connections should be opened 342. For example, parameters can be set for opening connections, e.g., based on the fee levied for using a particular connection, the connection speed, power drain (e.g., battery power required), and security. The system can scan available passive connections 230 and active connections 240 for a variety of criteria. For example, the mobile client 110 can assess the number of current connections (e.g. active or passive connections), their signal strength, and connection speed. The mobile client 110 can also detect its own geographic location to determine if it is in transit or leaving the effective range of a current connection, e.g., by using the global positioning system (GPS). Such parameters can be used to determine if additional connections should be

opened 342, if active connections 240 should be switched 344, and if existing connections should be closed.

Similarly, the mobile client 110 and/or home agent 160 can be programmed to determine if the active connection 240 should be switched 360. Parameters for the decision can include the relative signal strength, latency, transmittal rate, security, cost (e.g., charge per minute or per byte), power drain, and reliability of the different connections.

The process of switching 360 one connection from active 240 to passive 230 and another from passive 230 to active 240 is termed a "handoff." The use of common information exchange protocols for all connections can facilitate handoff. For example, TCP timers, packet size and windows do not need to be reconfigured for handoff. This is important for abrupt handoffs when the active connection 240 is unexpectedly lost and a passive connection 230 must be promoted to active 240.

Referring to Figure 4, an exemplary device 110 is equipped with a memory 412, and a processor 413 for executing instructions. The device 110 also has a console 414 (e.g., including a keyboard and display) for sending results to the user and for receiving user commands. The device 410 includes a chipset 420 or other interface for connecting the device's subsystems. The chipset 420 includes an input/output (I/O) interface 430 such as an I/O controller hub (ICH) or a basic I/O system (BIOS). For example, the ICH can be the Intel® 82801 ICH. The I/O interface is connected to two or more I/O units 432 and 434. These I/O units 432 and 434 are used to

establish communication channels with the home agent 160. For example, the I/O units can include a modem, (e.g., a modem connected to an internal or external mobile phone, a cable modem), an I/R port, a network card, BlueTooth, packet 5 cellular communicator, a satellite connection or a wireless interface (e.g., a wireless interface for CDPD, GPRS, GSM, 802.11a, 802.11b or BlueTooth).

Referring to Figure 5, information from an application 510 running on a mobile client 110 is mapped to an application 10 560 running on a destination server 170 (i.e. a server other than the home agent 160). The mobile client 110 is connected to the destination server 170 through the home agent 160. The information is passed to a rate limiter 520 that monitors 15 information transfer rates and quality of service parameters of the application.

In some examples, the rate limiter 520 implements the desired quality of service policy such that some application data is given bandwidth priority over other application data. The rate limiter 520 also guards against the mobile client 110 20 sending data to the home agent 160 faster than the home agent 160 can forward the data to its final destination, e.g., a destination server 170. The home agent 160 can instruct the rate limiter 520 on the client to slow down or stop sending 25 data for a specific destination (e.g., the destination server 170) when the home agent 160 is excessively buffering data. This situation can occur, for example, when the mobile client 160 is sending data to a system with a slower wireless

connection (e.g., the mobile client 160 is sending information to another mobile client through the home agent 160).

The flow indexer 521 monitors outgoing and incoming requests and information for the application. The flow indexer 521 can also frame information from the application with information for other applications. The mobility buffer 522 stores outgoing information (e.g., as packets) until an acknowledgement of its receipt is obtained from the home agent 160. If not receipt is received after an interval, the active connection 240 can be switched and the information (e.g., information packet) can be resent. Likewise, the mobility buffer 522 issues acknowledgments for information received from the home agent 160. The acknowledgements can be sent occasionally, e.g., along with outgoing data, to clear the other-side's mobility buffer. The acknowledgment can be time-insensitive. For example, the other-side's mobility buffer can send additional information before acknowledgments are received. The number of acknowledgments sent can be tailored to reflect the capacity of the other-size's mobility buffer. The mobility buffer sizes can be negotiated when opening the first connection between the client and home agent.

The information is then sent across a firewall 523. This process can include compression or decompression and/or encryption or decryption. Next, the information is sent to the transport socket 524 of the active connection 240.

The information is sent through the active connection 240 from the transport socket 524 on the mobile device 110 to the

transport socket 534 on the home agent 160. The home agent 160 processes the information in a corresponding manner with firewall traversal 533, mobility buffering and acknowledgement 532, application socket control 531, which similarly 5 disassembles framed information and routes the data to its final destination on the network (e.g., an intranet or the Internet), and mobility rate control 530. These processes map incoming information to an application socket 540 on the home agent 160. This socket, in turn, can be linked to an 10 application socket 550 on a destination server 170 that is networked to the home agent 160 by a VPN 150. Although the processes described above relates to information sent from the mobile device 110 to the home agent 160, the same processes can be executed for information flowing from the home agent 15 160 to the mobile device 110.

Application-specific information may be directed into the active connection 240 in various ways. Applications can be individually notified of the active connection 240 and can themselves direct information to the active connection 240. 20 In other examples, all application-specific information is directed to a network interceptor that routes the information to the active connection 240. For example, applications network calls and data may be captured directly from the applications and forwarded to the home agent 160. In other 25 examples, applications use local proxies, e.g., proxies running on the mobile client 110 to send traffic to the active connection 240.

In some examples, "sockets" are created as software objects which connect an application input/output stream to a network interface. Proxies are also used to process information from other applications.

5 Referring to Figure 6 and proceeding from right to left, the mobile device 110 is running applications that use three application sockets 610, 611, and 612. An application that requires a network connection can employ one or more of the application sockets 610, 611, and 612. Information exchange
10 to and from a socket is monitored by the flow indexer 616 and by socket control 620.

Outgoing information is processed by the flow indexer 616 and the rate limiter 622. Rate limiter 622 provides quality of service and protects against overflow of the home agent 160 buffer; the home agent 160 provides the rate limiter 622 with flow control notifications. Flow control systems (e.g., XON/XOFF or rate-based flow control systems) monitor information transfer rates, buffer load, and receipt acknowledgements. Flow control can be implemented on application and transport sockets on both the mobile client 110 and the home agent 160. For example, if the home agent 160 is unable to send information to a third destination 170 as fast as it receives it from the mobile client 110, the overflowing socket on the mobile client 110 is notified. Such measures can prevent the home agent 160 from buffering excessive data.
20
25

Outgoing information from multiple application sockets is packaged, compressed, encrypted and framed, by units including the upper encoder 624 and the lower encoder 630 via the mobility buffer 626. The mobility buffer 626 keeps a copy of the outgoing data until reception is acknowledged. If the active connection is lost and re-established (possibly on a different network interface), the content of mobility buffer 626 can be resent immediately, or, the computer (home agent or client) can wait for a mobility buffer acknowledgement (sent when a connection becomes active) before resending the exact amount of data the other side requires from the mobility buffer. The decision to send or wait before sending can be made based on the amount of data in the mobility buffer, the estimated bit rate of the active interface, cost of the interface and other values.

Router 640 directs traffic from the lower encoder 630 to the active transport socket (e.g., one of 660, 670, or 680). The active transport socket is a dedicated portal for information flow to the active connection 240. Information from the transport sockets to the application sockets is bi-directional. Information is secured using TLS or SSL protocols 666. The information is sent across the network firewall by the firewall traversal unit 664. Overflow information is buffered by the mobile client 110 using the outgoing overflow buffer 662. This information is stored until the home agent 160 acknowledges receipt of the contents, e.g., each segment of the contents).

Each transport socket 660, 670 and 680 is also monitored by a transport monitor 661, 671, and 681. General transport control 692 is overseen by the mobility control unit 642. General information flow is also monitored by the Statistics 5 (Stats) collector 646.

Incoming information is processed to traverse the firewall 664, and then verified and decrypted by the TLS/SSL unit 666. Once processed the information is sent via unit 640 to the lower decoder 636, the lower dispatcher 638 and the 10 upper decoder 634 and the upper dispatcher 632. These units can effect multiple information processing steps, e.g., removing framing information and decompressing and decrypting the information stream or packet. For example, socket control 620 routes information packets to the appropriate application 15 socket 610, 611 and 612.

The flow indexer 622 determines the information transmittal rate of the transport socket and appropriately buffers and/or compresses the information. For example, for a slow connection, information is compressed to minimize 20 information size, e.g., number of bytes, despite overhead in computational time to process the information. In contrast, for rapid connections, e.g., when the transmittal rate is not limiting, information is either not compressed or compressed using rapid algorithms that achieve more modest reduction in 25 information size.

Information transmittal rates can also monitored by applications to indicate the capacity of the active connection

240. For example, if capacity is available, applications can process low-level background networking tasks, such as backing up and updating files. If capacity is limited, applications can reduce their communications burden, e.g., by requesting 5 text or WAP pages instead of graphics.

Referring now to Figure 7 and proceeding from left to right, the home agent 160 has a process configuration similar to the one for the mobile client 110. Information from connections is mapped to transport sockets 760 and 770. The 10 information is sent for firewall traversal 764 and 774 and then for security verification and decryption 766 and 776. A router 740 directs incoming information to the lower decoder 736. Outgoing information is also directed by the router 740 toward the transport sockets 760 and 770 by way of output 15 overflow buffers 762 and 772.

Mobility control 742 monitors the information overflow in these buffers 762 and 772, and can appropriately regulate outgoing information through the rate limiter 772, the router 740, and the lower encoder 730.

20 Incoming information is processed by the lower decoder 736 and the upper decoder 734. These decoders are linked to the lower dispatcher 738, which can issue tunnel or connection specific commands, and the upper dispatcher 732, which can issue application specific commands. The decoded information 25 is dissembled into individual items and then distributed to the correct application sockets 710, 711, and 712 on the home agent 160. Hence, information from the application sockets

610, 611, and 612 running on the home agent 160, is mapped to corresponding application sockets 710, 711, and 712 on the home agent 160. As described above, these application sockets also generate outgoing information for delivery to the mobile device 110.

5 The information can be transmitted through the active connection 240 by a variety of methods, for example, in packets or as a continuous stream.

The information sent as packets can be framed. The 10 information from different application sockets can be aggregated such that individual commands and data items from different applications are included in a single information packet. This packet is encrypted and compressed as a single unit using the previously initialized protocols thus obviating 15 the need to include authentication, cryptographic, and compression headers. The packet is framed with minimal header information, such as a checksum and a packet identifier. On reception of the packet, an acknowledgement is issued for the entire packet, e.g., an identifier and/or checksum for the 20 received information. The aggregation of data is in contrast to systems that emit individual data packets for each network call made or data item sent by individual applications. In addition to reducing overhead, framing multiple items together enables larger individual information packets, which can 25 include related items. By placing related items together, the latency due to packet mis-alignment, such as occurs when conventional packets arrive out of order, is reduced.

Optionally, the protocol includes the reduction of "round trip" information exchanges. When desired, this feature determines if multiple information exchanges between the mobile device and a destination system 170 (i.e., a system other than the home agent 160) can be emulated by a proxy module running on the home agent 160. The proxy module interacts with the destination system 170 on behalf of the mobile client 110 without having to exchange information with the mobile client 110. The use of the proxy module reduces network traffic and overcomes inefficiencies due to the latency of some wireless connections. Proxy modules can be generic, e.g., able to support more than one application, or can be dedicated, e.g., application specific.

For example, the use of a web browser requires at least three round trips that can be obviated by a proxy module. Conventionally, when a user enters a uniform resource locator (URL) for a web site, the web browser on the mobile client 110 first sends a request to a domain name server to obtain the actual IP address of the URL. The mobile client 110 then initiates a TCP connection with the web server at the IP address. Only after confirmation does the mobile client 110 send a hypertext transfer protocol (HTTP) request to the web server and subsequently receive the hypertext content.

The use of a proxy module for the web browser instead replaces the three network roundtrips with a single exchange between the home agent 160 and the mobile client 110, thus reducing information traffic in the active connection 240.

When a user enters a URL for a web site, this URL request is sent on the active connection 240 to a proxy module on the home agent 160 or on a destination server 170. The proxy module queries the domain name server, initiates the TCP connection, and sends the HTTP request. These steps are done without contacting the mobile client 110. When the proxy module receives the hypertext content, it directs it back over the active connection 240 to the mobile client 110.

An additional method to reduce network traffic is the use of application specific content re-purposing. This feature reformats information, particularly content. Extraneous content can be removed to reduce the amount of data. For example, data is translated by the home agent 160 from hypertext markup language (HTML) and graphics formats to i-mode or WAP compatible information.

Frequently, the client 110 and/or server 160 have more than one application sending information along a single active connection. Quality of service protocols are utilized to determine which application has priority for communicating information. Quality of service parameters can be implemented at many levels, for example, at the level of the upper encoder 624 and lower encoder 630 which frame network calls and data for different applications. The relative content of each framed data packet can be varied in accordance to quality of service parameters.

Further, the home agent 160, as well as the mobile client 110, can determine quality of service. Quality of service can

be asymmetric such that an application may have high priority on the home agent 160 to send data to the mobile device 110, but may have low priority on the mobile device 110 for sending data to the home agent 160. A quality of service parameter 5 indicating priority for information exchange can be transmitted from the home agent 160 to the mobile device 110 or from the mobile device 110 to the home agent 160.

Other implementations are within the scope of the claims.

For example, the techniques described here are not 10 limited to any particular hardware or software configuration; they may find applicability in any computing or processing environment. The techniques may be implemented in hardware, software, or a combination of the two. The techniques may be implemented in programs executing on programmable machines 15 such as mobile or stationary computers, personal digital assistants, and similar devices that each include a processor, a storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and one or more output devices.

Each program may be implemented in a high level 20 procedural or object oriented programming language to communicate with a machine system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted 25 language.

Each such program may be stored on a storage medium or device, e.g., compact disc read only memory (CD-ROM), hard

disk, magnetic diskette, or similar medium or device, that is readable by a general or special purpose programmable machine for configuring and operating the machine when the storage medium or device is read by the computer to perform the 5 procedures described in this document. The system may also be implemented as a machine-readable storage medium, configured with a program, where the storage medium so configured causes a machine to operate in a specific and predefined manner.

The processes described here may be executed by an 10 embedded system.